PTO/SB/17 (12-04)
Approved for use through 07/31/2006. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Effective on 12/08/2004.
Fees pursuant to the Consolidated Appropriations Act, 2005 (H.R. 4818).

FEE TRANSMITTAL
For FY 2005

☐ Applicant claims small entity status. See 37 CFR 1.27

TOTAL AMOUNT OF PAYMENT ($) 500.00

Complete if Known	
Application Number	09/778,300
Filing Date	February 7, 2002
First Named Inventor	Marc Adam Kaplan
Examiner Name	Ryman, Daniel J.
Art Unit	2665
Attorney Docket No.	YOR920000846US1 (8728-472)

METHOD OF PAYMENT (check all that apply)

☐ Check ☐ Credit Card ☐ Money Order ☐ None ☐ Other (please identify): _____

☒ Deposit Account Deposit Account Number: ~~50-0150~~ 50-0510 Deposit Account Name: IBM/YORKTOWN HEIGHTS

For the above-identified deposit account, the Director is hereby authorized to: (check all that apply)

☒ Charge fee(s) indicated below ☐ Charge fee(s) indicated below, except for the filing fee

☒ Charge any additional fee(s) or underpayments of fee(s) under 37 CFR 1.16 and 1.17 ☐ Credit any overpayments

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

FEE CALCULATION

1. BASIC FILING, SEARCH, AND EXAMINATION FEES

	FILING FEES		SEARCH FEES		EXAMINATION FEES		
		Small Entity		Small Entity		Small Entity	
Application Type	Fee ($)	Fee ($)	Fee ($)	Fee ($)	Fee ($)	Fee ($)	Fees Paid ($)
Utility	300	150	500	250	200	100	
Design	200	100	100	50	130	65	
Plant	200	100	300	150	160	80	
Reissue	300	150	500	250	600	300	
Provisional	200	100	0	0	0	0	

2. EXCESS CLAIM FEES

Fee Description	Fee ($)	Small Entity Fee ($)
Each claim over 20 or, for Reissues, each claim over 20 and more than in the original patent	50	25
Each independent claim over 3 or, for Reissues, each independent claim more than in the original patent	200	100
Multiple dependent claims	360	180

Total Claims	Extra Claims	Fee ($)	Fee Paid ($)	Multiple Dependent Claims Fee ($)	Multiple Dependent Claims Fee Paid ($)
_____ - 20 or HP =	_____ x	_____ =	_____	_____	_____

HP = highest number of total claims paid for, if greater than 20

Indep. Claims	Extra Claims	Fee ($)	Fee Paid ($)
_____ - 3 or HP =	_____ x	_____ =	_____

HP = highest number of independent claims paid for, if greater than 3

3. APPLICATION SIZE FEE

If the specification and drawings exceed 100 sheets of paper, the application size fee due is $250 ($125 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).

Total Sheets	Extra Sheets	Number of each additional 50 or fraction thereof	Fee ($)	Fee Paid ($)
_____ - 100 =	_____ / 50 =	_____ (round up to a whole number) x	_____ =	_____

4. OTHER FEE(S)

	Fees Paid ($)
Non-English Specification, $130 fee (no small entity discount)	
Other: Appeal Brief	500.00

SUBMITTED BY

Signature		Registration No. (Attorney/Agent) 48,909	Telephone 516-692-8888
Name (Print/Type)	Nathaniel T. Wallace		Date October 19, 2005

AF ✓ JPW

PTO/SB/21 (09-04)
Approved for use through 07/31/2006. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

TRANSMITTAL FORM

(to be used for all correspondence after initial filing)

Total Number of Pages in This Submission

Application Number	09/778,300
Filing Date	February 7, 2002
First Named Inventor	Marc Adam Kaplan
Art Unit	2665
Examiner Name	Ryman, Daniel J.
Attorney Docket Number	YOR920000846US1 (8728-472)

ENCLOSURES *(Check all that apply)*

- ☐ Fee Transmittal Form
  - ☐ Fee Attached
- ☐ Amendment/Reply
  - ☐ After Final
  - ☐ Affidavits/declaration(s)
- ☐ Extension of Time Request
- ☐ Express Abandonment Request
- ☐ Information Disclosure Statement
- ☐ Certified Copy of Priority Document(s)
- ☐ Reply to Missing Parts/ Incomplete Application
  - ☐ Reply to Missing Parts under 37 CFR 1.52 or 1.53
- ☐ Drawing(s)
- ☐ Licensing-related Papers
- ☐ Petition
- ☐ Petition to Convert to a Provisional Application
- ☐ Power of Attorney, Revocation
- ☐ Change of Correspondence Address
- ☐ Terminal Disclaimer
- ☐ Request for Refund
- ☐ CD, Number of CD(s) _____
  - ☐ Landscape Table on CD
- ☐ After Allowance Communication to TC
- ☐ Appeal Communication to Board of Appeals and Interferences
- ☒ Appeal Communication to TC (Appeal Notice, Brief, Reply Brief)
- ☐ Proprietary Information
- ☐ Status Letter
- ☐ Other Enclosure(s) (please identify below):

Remarks

SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT

Firm Name	F. Chau & Associates, LLC		
Signature			
Printed name	Nathaniel T. Wallace		
Date	October 19, 2005	Reg. No.	48,909

CERTIFICATE OF TRANSMISSION/MAILING

I hereby certify that this correspondence is being facsimile transmitted to the USPTO or deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on the date shown below:

Signature			
Typed or printed name	Nathaniel T. Wallace	Date	48,909

PATENTS

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

APPLICANT(S):	Kaplan	DOCKET: YOR920000846US1 (8728-472)
SERIAL NO.:	09/778,300	GROUP ART UNIT: 2665
FILED:	February 7, 2001	EXAMINER: Ryman, Daniel J.
FOR:	**SYSTEM AND METHOD FOR A MULTICAST NETWORK MESSAGING SERVICE**	

Mail Stop Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

RESPONSE TO NOTICE OF NON-COMPLIANCE WITH APPEAL BRIEF

In response to the Advisory Action dated August 5, 2005 and the Final Office Action dated May 17, 2005, finally rejecting Claims 1-14 and 23-28 under 35 U.S.C. §103(a). Applicant appeals pursuant to the Notice of Appeal filed on August 17, 2005 and submit this appeal brief.

11/03/2005 JBALINAN 00000004 500510 09778300
01 FC:1402 500.00 DA

CERTIFICATE OF MAILING UNDER 37 C.F.R. §1.8(a)

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail, postpaid in an envelope, addressed to the: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on October 19, 2005.

Dated: October 19, 2005

Melithza Rodriguez

TABLE OF CONTENTS

1. Real Party in Interest

The real party in interest is International Business Machines Corporation, the assignee of the entire right, title, and interest in and to the subject application by virtue of an assignment of record.

2. Related Appeals and Interferences

None.

3. Status of Claims

Claims 1-14 and 23-28 are pending, stand rejected, and are under appeal.

Claims 15-22 have been cancelled.

A copy of the Claims as pending is presented in the Appendix.

4. Status of Amendments

Claims 1, 4 and 14 were amended by the Amendment under 37 C.F.R. §1.111, filed February 28, 2005. By the Amendment claims 15-22 were cancelled and claims 23-28 were added. This Amendment was entered.

An Amendment was submitted under 37 C.F.R. §1.116, filed July 14, 2005 to amend Claims 1, 5, 6 and 14, and cancel Claim 4. By the Advisory Action dated August 5, 2005, this Amendment was not entered.

5. Summary of Claimed Subject Matter

The present invention relates to distribution of messages within a network of processors, and more particularly, to a high capacity, fault-tolerant multicast message distribution system.

A multicast spanning tree is defined over a set of cells and link bundles, wherein each cell includes at least one message processing computer and each link bundle includes at least one communication link.

Each processor can directly communicate with any other processor in the same cell. Moreover, any processor within a cell can directly multicast a message to any subset of processors within the same cell. A single message processor can be a member of more than one cell.

Referring to Claims 1 and 14, a method is claimed for the multicast distribution of a message from a first real machine through a network of message processing machines to one or more message receiving machines, wherein the network is organized into two or more cells including machines, and wherein one or more links between cells comprise link bundles (see Page 6, lines 9-12). The method includes selecting a spanning tree rooted in the cell containing the first real machine, and comprised of the cells and the link bundles (see Page 17, lines 8-11), and determining one or more receiving cells including a message receiving machine on the selected spanning tree (see

Page 24, lines 9-15). The method includes selecting a cellule in each of the one or more receiving cells to receive the message, wherein each cellule comprises a set of one or more virtual machines within a cell at an end of a link bundle (see page 24, lines 15-18), and selecting a route to each selected cellule from a routing choice table of the first real machine including potential routing choices to reachable cellules relative to the first real machine (see Page 24, lines 18-19). The method includes routing the message to each selected cellule in the one or more receiving cells in the spanning tree, and delivering the message to each message receiving machine within the one or more receiving cells (see Page 24, line 22 to Page 23, line 7). Particular to Claim 14, a program storage device is described at, Page 6, lines 16 to Page 7, line 7).

Referring to Claim 23, a method is claimed for the multicast distribution of a message from a publishing client through a network of message processing machines to a first subscribing client, wherein the network is organized into two or more cells including client machines, and wherein one or more links between cells comprise link bundles (see Page 6, lines 9-12). The method includes determining, by the first subscribing client, whether the publishing client is an external neighbor outside a first cell of the first subscribing client or an internal neighbor inside the first cell of the first subscribing

client (see Page 24, lines 10-12), forwarding the message from the first subscribing client to an internal subscribing neighbor client upon determining the publishing client to be an external neighbor (see Page 18, lines 5-10), and forwarding the message from the first subscribing client to an external subscribing neighbor client upon determining the publishing client to be an internal neighbor (see Page 18, lines 11-15).

The routing methods facilitate load balancing by having each message processor dynamically choose among the different virtual machines. Therefore, incoming messages bound towards a given cellule are distributed among the multiple message processors that implement the virtual machines of that cellule, exploiting either multiple inter-cellular network connections or intracellular networking facilities.

6. Grounds of Rejection to be Reviewed on Appeal

A. Claims 1-3, 14 and 14 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Doeringer et al. (U.S. Patent No. 5,361,256) in view of Cheng (U.S. Patent No. 6,600,724) and further in view of Coile et al. (U.S. Patent No. 6,061,349).

B. Claims 4-12 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Doeringer et al. (U.S. Patent No. 5,361,256) in view of Cheng (U.S. Patent No. 6,600,724) in view of Coile et al. (U.S. Patent No. 6,061,349) and further in view of Annapareddy et al. (U.S. Patent No. 5,602,839).

C. Claims 23, 25 and 28 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Doeringer et al. (U.S. Patent No. 5,361,256) in view of Cheng (U.S. Patent No. 6,600,724).

D. Claims 24, 26 and 27 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Doeringer et al. (U.S. Patent No. 5,361,256) in view of Cheng (U.S. Patent No. 6,600,724) and further in view of Coile et al. (U.S. Patent No. 6,061,349).

7. Argument

A. The Claim Rejections Under 35 U.S.C. 103 Are Legally Deficient.

In rejecting claims under 35 U.S.C. §103, the Examiner bears the initial burden of presenting a *prima facie* case of obviousness. In re Rijckaert, 9 F.3d 1531, 1532 (Fed. Cir. 1993). The burden of presenting a *prima facie* case of obviousness is only satisfied by showing some objective teaching in the prior art or that knowledge generally available to one of ordinary skill in the art would lead that individual to combine the relevant teachings of the references. In re Fine, 837 F.2d 1071, 1074 (Fed. Cir. 1988). A *prima facie* case of obviousness is established when the teachings of the prior art itself would appear to have suggested the claimed subject matter to one of ordinary skill in the art. In re Bell, 991 F.2d 781, 782 (Fed. Cir. 1993). If the Examiner fails to establish a *prima facie* case, the rejection is improper and must be overturned. In re Rijckaert, 9 F.3d at 1532 (citing In re Fine, 837 F.2d at 1074).

i. Claims 1-3, 13 and 14

It is respectfully submitted that at the very least, the combined teachings of Doeringer, Cheng and Coile are legally

deficient to establish a *prima facie* case of obviousness against independent Claims 1 and 14.

Claims 1 and 14 claim, *inter alia*, "selecting a route to each selected cellule from a routing choice table of the first real machine including potential routing choices to reachable cellules relative to the first real machine."

The combined teachings of Doeringer, Cheng and Coile are legally deficient to establish a prima facie case of obviousness against Claims 1 and 14 because Doeringer, Cheng and Coile do not teach or suggest "selecting a route to each selected cellule from a routing choice table of the first real machine including potential routing choices to reachable cellules relative to the first real machine."

Doeringer teaches a method for multicast routing, and sending a multicast packet to a subnetwork (see col. 10, line 20 to col. 11, line 3). Doeringer does not teach "selecting a route to each selected cellule from a routing choice table of the first real machine including potential routing choices to reachable cellules relative to the first real machine" as claimed in claims 1 and 14. Doeringer routes multicast packets according to end systems having a common groupid within one or more subnetworks (see col. 7, lines 54-61 and col. 10, line 20 to col. 11, line 3). Doeringer's network uses multicast routing tables (MURT) maintained at gateways (see col. 10, lines 24-35).

Doeringer teaches that a source node of the multicast packets does not select a route to a selected cellule from a routing table. The source node sends a multicast, having been addressed to a groupid, to a gateway (see col. 10, line 23-24). Doeringer is silent about the source's sending of the multicast; Doeringer does not teach that the sending includes a selection from a routing choice table. The gateway of Doeringer does select a route to a next hop (see col. 10, lines 27-29 and lines 48-50). However, the gateway is not a first real machine as claimed in Claims 1 and 14, for example, a spanning tree is not rooted in the gateway. Therefore, Doeringer fails to teach "selecting a route to each selected cellule from a routing choice table of the first real machine," essentially as claimed in Claims 1 and 14.

Cheng teaches a routing scheme for a shortest path tree architecture (see Abstract). Cheng does not teach or suggest "selecting a route to each selected cellule from a routing choice table of the first real machine including potential routing choices to reachable cellules relative to the first real machine" as claimed in Claims 1 and 14. Cheng does not teach or suggest selecting a route to each selected cellule from a routing choice table of a first real including potential routing choices, essentially as claimed in Claims 1 and 14. Cheng's routing scheme does not implement a routing choice table (see

col. 7, lines 10-18). Cheng teaches that a route is given by an address and a shortest path tree containing the address (see col. 7, lines 13-15). The route of Cheng is not selected from a routing choice table. Therefore, Cheng fails to teach or suggest "selecting a route to each selected cellule from a routing choice table of the first real machine including potential routing choices to reachable cellules relative to the first real machine" as claimed in Claims 1 and 14. Thus, Cheng fails to cure the deficiencies of Doeringer.

Coile teaches a system and method for handling a plurality of connection requests made for a plurality of virtual machines implemented on a single physical machine (see col. 2, lines 36-43). Coile does not teach or suggest "selecting a route to each selected cellule from a routing choice table of the first real machine including potential routing choices to reachable cellules relative to the first real machine" as claimed in Claims 1 and 14. Coile's virtual machines each have an individual IP address, which are mapped to an IP address of a physical machine by a Local Director; the physical machine supporting multiple connections to clients on a single network communication line (see col. 6, lines 43-61 and col. 4, lines 60-67). Clients of Coile's virtual machines communicate using an IP address. Knowledge of a destination's IP address is not analogous to selecting a route, essentially as claimed in claims

1 and 14. An IP address specifies a destination without regard for any particular route. Further, nowhere does Coile teach that an IP address is selected from a routing table. Therefore, Coile does not teach or suggest "selecting a route to each selected cellule from a routing choice table of the first real machine including potential routing choices to reachable cellules relative to the first real machine" as claimed in Claims 1 and 14. Thus, Coile fails to cure the deficiencies of Doeringer and Cheng.

It is respectfully submitted that the combined teachings of Doeringer, Cheng and Coile fail to teach or suggest "selecting a route to each selected cellule from a routing choice table of the first real machine including potential routing choices to reachable cellules relative to the first real machine" as claimed in Claims 1 and 14.

Accordingly, the rejection of Claims 1-3, 13 and 14 should be overruled. Claims 2, 3 and 13 depend from Claim 1. The dependent claims include the elements of independent Claim 1 and they are not rendered obvious by the cited references for at least the reasons given for the independent Claim 1.

ii. Claims 4-12

Claims 4-12 depend from claim 1. The dependent claims include the elements of independent Claim 1 and they are not

rendered obvious by the cited references for at least the reasons given for the independent Claim 1.

iii. Claims 23, 25 and 28

It is respectfully submitted that at the very least, the combined teachings of Doeringer and Cheng are legally deficient to establish a *prima facie* case of obviousness against independent Claim 23.

Claim 23 claims "determining, by the first subscribing client, whether the publishing client is an external neighbor outside a first cell of the first subscribing client or an internal neighbor inside the first cell of the first subscribing client."

Doeringer teaches a method for multicast routing, and sending a multicast packet to a subnetwork (see col. 10, line 20 to col. 11, line 3). Doeringer does not teach "determining, by the first subscribing client, whether the publishing client is an external neighbor outside a first cell of the first subscribing client or an internal neighbor inside the first cell of the first subscribing client" as claimed in Claim 23. Doeringer teaches that multicast packets are sent to gateways with a destination field and a target subnetwork field (see col. 10, lines 30-34). Doeringer performs messaging based on destinations and targets. The gateways of Doeringer are aware of

destinations, but do not have knowledge of senders, much less any concept of a cell having an interior and exterior. For example, as shown in Figure 5, gateways are stand-alone devices, separate from subnets and are not part of any cell. Therefore, Doeringer fails to teach "determining, by the first subscribing client, whether the publishing client is an external neighbor outside a first cell of the first subscribing client or an internal neighbor inside the first cell of the first subscribing client", essentially as claimed in Claim 23. Therefore, Doeringer fails to teach all the limitations of Claim 23.

Cheng teaches a routing scheme for a shortest path tree architecture (see Abstract). Cheng does not teach or suggest "determining, by the first subscribing client, whether the publishing client is an external neighbor outside a first cell of the first subscribing client or an internal neighbor inside the first cell of the first subscribing client", essentially as claimed in Claim 23. Cheng's routing scheme consults an address table to determine a destination node (see col. 6, lines 61-65). Cheng teaches routing according to a location of a destination and a shortest path tree (see col. 7, lines 13-15). The address tables of Cheng include information about destinations relative to an intermediary (e.g., switch) between a source and destination. The tables do not include information about the source. Therefore, Cheng fails to teach or suggest,

"determining, by the first subscribing client, whether the publishing client is an external neighbor outside a first cell of the first subscribing client or an internal neighbor inside the first cell of the first subscribing client", essentially as claimed in Claim 23. Thus, Cheng fails to cure the deficiencies of Doeringer.

The combined teachings of Doeringer and Cheng teach methods for determining where to route a message or packet according to a destination. The combined teachings of Doeringer and Cheng fail to teach or suggest a determination of a location of a publishing client. Therefore, Claim 23 is believed to be allowable in view of Doeringer and Cheng. Claims 25 and 28 depend from Claim 23 and are believed to be allowable for at least the reasons given for Claim 23.

iv. Claims 24, 26 and 27

Claims 24, 26 and 27 depend from Claim 23. The dependent claims are believed to be allowable for at least the reasons given for Claim 23. The Examiner's reconsideration of the rejection is respectfully requested.

C. CONCLUSION

The claimed invention is not disclosed or suggested by the teachings of the applied prior art references, either alone or in combination. Moreover, the Examiner has failed to establish a *prima facie* case of obviousness of the presently claimed methods under 35 U.S.C. §103 over Doeringer, Cheng and Coile with respect to Claims 1, 14, and under 35 U.S.C. §103 over Doeringer and Cheng with respect to Claim 23 for at least the reasons noted above. Claims 2-14 depend from Claim 1. Claims 24-28 depend from Claim 28. The dependent claims are believed to be allowable for at least the reasons given for Claims 1 and 23. Accordingly, it is respectfully requested that the Board overrule the rejection of Claims 1–14 and 23-28 under 35 U.S.C. §103.

Date: October 19, 2005

By: 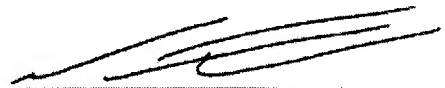

Nathaniel T. Wallace
Reg. No. 48,909
Attorney for Appellants

F. CHAU & ASSOCIATES, LLP
130 Woodbury Road
Woodbury, New York 11797
TEL: (516) 692-8888
FAX: (516) 692-8889

8. CLAIMS APPENDIX

What is claimed is:

1. A method for the multicast distribution of a message from a first real machine through a network of message processing machines to one or more message receiving machines, wherein the network is organized into two or more cells including machines, and wherein one or more links between cells comprise link bundles, the method comprising the steps of:

selecting a spanning tree rooted in the cell containing the first real machine, and comprised of the cells and the link bundles;

determining one or more receiving cells including a message receiving machine on the selected spanning tree;

selecting a cellule in each of the one or more receiving cells to receive the message, wherein each cellule comprises a set of one or more virtual machines within a cell at an end of a link bundle;

selecting a route to each selected cellule from a routing choice table of the first real machine including potential routing choices to reachable cellules relative to the first real machine;

routing the message to each selected cellule in the one or more receiving cells in the spanning tree; and

delivering the message to each message receiving machine within the one or more receiving cells.

2. The method of claim 1, further comprising the step of implementing one or more virtual machines within a real machine.

3. The method of claim 1, wherein a link is one of a virtual link between two virtual machines, and a real link between two real machines.

4. The method of claim 1, wherein the multicast distribution of the message is along links and further comprises the step of routing the message through the selected spanning tree according to precomputed cellule distribution tables associated with the each virtual machine, wherein each cellule distribution table includes a first distribution set of cellules to be used if the message is received from a neighbor in the same cell and a second distribution set of cellules to be used if the message is received from a neighbor in another cell.

5. The method of claim 4, wherein the step routing further comprises the step of determining a routing choice table for each real machine.

6. The method of claim 4, wherein the multicast distribution is according to the cellule distribution table and a message distribution tag including a flagged list of virtual machines.

7. The method of claim 5, wherein the routing choice table selects machines and links according to one of random choice, round-robin least busy, least-busy, preserve message order, and preserve message order by hashing on origin identification.

8. The method of claim 5, wherein the step of determining a routing choice table further includes the step of determining a failover route for redirecting a message.

9. The method of claim 5, wherein the step of determining a routing choice table further includes the step of exchanging routing information included in the routing choice table of each machine upon the happening of an event.

10. The method of claim 9, wherein an event includes one of a machine failure and a machine recovery.

11. The method of claim 6, wherein the message distribution tags can be one of compressed, factored between internal and external machines relevant to a sending machine, and compressed and

factored.

12. The method of claim 11, further comprising the step of determining an updated message distribution tag for the message relevant to the internal and external machines of the sending machine, wherein the sending machine can be one of the first real machine and a receiving machine for forwarding the message to one or more additional receiving machines.

13. The method of claim 1, further comprising the step of scaling the message handling capacity of the network.

14. A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for the multicast distribution of a message from a first real machine through a network of message processing machines to one or more message receiving machines, wherein the network is organized into two or more cells including machines, and wherein one or more links between cells comprise link bundles, the method steps comprising:

selecting a spanning tree rooted in the cell containing the first real machine, and comprised of the cells and the link bundles;

determining one or more receiving cells including a message

receiving machine on the selected spanning tree;

selecting a cellule in each of the one or more receiving cells to receive the message, wherein each cellule comprises a set of one or more virtual machines within a cell at an end of a link bundle;

selecting a route to each selected cellule from a routing choice table of the first real machine including potential routing choices to reachable cellules relative to the first real machine;

routing the message to each selected cellule in the one or more receiving cells in the spanning tree; and

delivering the message to each message receiving machine within the one or more receiving cells.

23. A method for the multicast distribution of a message from a publishing client through a network of message processing machines to a first subscribing client, wherein the network is organized into two or more cells including client machines, and wherein one or more links between cells comprise link bundles, the method comprising:

determining, by the first subscribing client, whether the publishing client is an external neighbor outside a first cell of the first subscribing client or an internal neighbor inside the first cell of the first subscribing client;

forwarding the message from the first subscribing client to an internal subscribing neighbor client upon determining the publishing client to be an external neighbor; and

forwarding the message from the first subscribing client to an external subscribing neighbor client upon determining the publishing client to be an internal neighbor.

24. The method of claim 23, further comprising excluding a subscribing client in a cellule in which the first subscribing client implements a virtual machine, wherein each cellule is a disjoint subset of virtual machines within a cell.

25. The method of claim 23, further comprising excluding a subscribing external neighbor client in the first cell of the first subscribing client upon determining the publishing client to be an external neighbor.

26. The method of claim 23, further comprising excluding a virtual machine of a subscribing client within the first cell of the first subscribing client upon determining the publishing client to be an internal neighbor.

27. The method of claim 23, further comprising:

determining whether a cellule in a second cell of the first

subscribing client in which the first subscribing client implements a virtual machine has received the message via a virtual machine of a second subscribing client; and

excluding a subscribing client within the second cell from receiving the message from the first subscribing client.

28. The method of claim 23, further comprising tagging the message, wherein a tag indicates an intended recipient client, whether the message has been received from an internal neighbor or external neighbor, whether the message has been received from an internal neighbor virtual machine or from an external neighbor virtual machine.